

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

Multi-Channel Memory Architecture

Inventors:

Belgacem Haba

Sayeh Khalili

Donald R. Mullen

Nader Gamini

ATTORNEY'S DOCKET NO. RB1-008US

1 **TECHNICAL FIELD**

2 The present invention relates to memory systems and, in particular, to a
3 multi-channel memory architecture that provides increased communication
4 bandwidth and provides communication channels having conductors of
5 substantially equal length.
6

7 **BACKGROUND**

8 In recent years, DRAM (dynamic random access memory) technology has
9 progressed dramatically. Device storage capacities have increased from 1 Kbits
10 per chip to 64 Mbits per chip, a factor of 64,000. However, DRAM performance
11 has not maintained pace with these density changes, since access times have
12 decreased by a factor of approximately five. During this same time period,
13 microprocessor performance has increased by several orders of magnitude. This
14 disparity between the speed of the microprocessors and the access time associated
15 with DRAMs has forced system designers to create a variety of expensive and
16 often complicated hierarchical memory techniques, such as SRAM (static random
17 access memory) caches and parallel arrays of DRAMs. Some high performance
18 graphics systems rely on expensive frame buffers to provide the necessary
19 bandwidth for communicating with the memory devices.

20 Rambus Inc. has developed a new chip-to-chip bus, referred to as the
21 "Direct Rambus Channel", that operates up to ten times faster than conventional
22 DRAMs. The Direct Rambus Channel (the channel) connects memory devices to
23 other devices such as microprocessors, digital signal processors, graphics
24 processors and ASICs (application-specific integrated circuits). The channel uses
25

1 a small number of very high speed signals to carry all address, data, and control
2 information.

3 Fig. 1 illustrates a known memory module 100 containing multiple memory
4 devices and a single channel coupling the memory devices to one another.
5 Memory module 100 typically includes a substrate 102, such as a multi-layered
6 printed circuit board (PCB) that supports multiple memory devices 104 mounted
7 to the substrate. In this example, each memory device 104 is a Rambus DRAM
8 (or "RDRAM[®]"), developed by Rambus Inc. of Mountain View, California. The
9 memory devices are coupled to one another by a Direct Rambus Channel 106.
10 Data flows into channel 106 on the left side of module 100, as oriented in Fig. 1.
11 Data flows along channel 106, past the four RDRAMs 104, and out the right side
12 of module 100. In a typical configuration, one or more modules 100 are supported
13 by a motherboard (not shown), which includes a memory controller. Each of the
14 four RDRAMs 104 can retrieve data from channel 106 and transmit data to other
15 devices using channel 106.

16 Channel 106 shown in Fig. 1 provides a high bandwidth communication
17 path for address, data, and control information associated with one or more of the
18 four RDRAMs 104. Channel 106 includes multiple conductors, two of which are
19 identified by the reference numbers 106A and 106B. As shown in Fig. 1, the
20 channel conductors follow a ninety degree bend at each end of the substrate 102.
21 Thus, the lengths of the various conductors in channel 106 are not equal. For
22 example, the length of conductor 106A is greater than the length of conductor
23 106B in the same channel. This unequal routing of conductors in channel 106
24 causes problems with the timing of signals propagating on the channel because
25 different signals arrive at memory device 104(1) at different times due to the

unequal conductor lengths. Similarly, the conductors between memory device 104(4) and the channel output are unequal in length, leading to timing problems with signals on the channel output. To compensate for the unequal channel conductor lengths, the conductors may be routed on substrate 102 in such a manner that each conductor is approximately the same length. Attempting to route conductors such that each conductor is approximately the same length is tedious and time-consuming. Additionally, such routing of conductors increases the cost and complexity of the substrate 102 (e.g., by requiring additional layers of conductors and additional vias between layers).

Alternatively, if the channel conductors are maintained at unequal lengths, then the processing of signals from the channel must be delayed for a period of time that allows the signal on the longest conductor to reach the appropriate device (e.g., memory device). This alternative slows the operation of the overall memory system by causing an added delay to compensate for the worst case delay on the channel.

The memory module 100 shown in Fig. 1 is also limited to a single channel 106 on each side of the substrate 102 (i.e., a maximum of two channels per memory module). Thus, all memory devices on one side of the substrate 102 share a common channel, which is limited to a particular bandwidth. The speed at which data can be stored to or retrieved from the memory devices 104 is limited by the bandwidth of the single channel.

The memory architecture described herein addresses these and other problems by providing an architecture in which all channel conductors are substantially the same length. Further, the memory architecture described herein

1 supports the use of multiple channels on each side of a substrate, thereby
2 increasing the overall bandwidth.

3 4 5 SUMMARY

6 In one embodiment, a memory architecture includes a substrate having a
7 first elongated edge and a second elongated edge in which the two elongated edges
8 are opposite one another. Multiple memory devices are disposed on the substrate.
9 Multiple channels extend from the first elongated edge to the second elongated
10 edge such that each of the multiple memory devices is coupled to one of the
11 multiple channels

12 In another embodiment, multiple channels extend across both sides of the
13 substrate.

14 In a particular implementation, each channel includes multiple conductors
15 that follow a substantially linear path across the substrate.

16 In a described embodiment, each channel includes multiple conductors that
17 are approximately equal in length.

18
19 Another embodiment provides a first substrate containing multiple memory
20 devices and a first channel portion extending across the first substrate. A second
21 substrate also contains multiple memory devices and a second channel portion
22 extending across the second substrate. A connector is configured to couple the
23 first channel portion to the second channel portion. The connector includes a first
24 slot that receives an edge of the first substrate and a second slot that receives an
25 edge of the second substrate.

1 In a particular embodiment, the coupling of the first channel portion to the
2 second channel portion through the connector forms a channel
3
4

5 **BRIEF DESCRIPTION OF THE DRAWINGS**

6 Fig. 1 illustrates a known memory module containing multiple memory
7 devices and a single channel coupling the memory devices to one another.
8

9 Fig. 2 illustrates another memory module having multiple channels in
10 which all channel conductors are substantially the same length.
11

12 Fig. 3 illustrates a perspective view of a memory architecture including a
13 pair of memory modules and a pair of connectors to couple the memory modules
14 to one another.
15

16 Fig. 4 illustrates an exploded view of the memory architecture shown in
17 Fig. 3.
18

19 Fig. 5 is a top view of a connector used to couple two memory modules to
20 one another.
21

22 Fig. 6 is a flow diagram illustrating a procedure for creating and using a
23 memory architecture of the type described herein.
24

25 Fig. 7 illustrates another embodiment of a memory architecture having
multiple channels in which all channel conductors are substantially the same
length.

DETAILED DESCRIPTION

The multi-channel memory architecture described herein provides increased bandwidth by providing additional communication channels on each memory module. These additional channels increase the overall bandwidth of the memory module. Additionally, the described memory architecture routes the channel conductors such that each channel has conductors of substantially equal length. This routing of the channel conductors simplifies the layout of the conductors on the substrate and eliminates the need for special timing consideration caused by channel conductors of unequal length.

Fig. 2 illustrates another memory module 200 having multiple channels in which all channel conductors are substantially the same length. Memory module 200 eliminates the ninety degree bend in the channel conductors, thereby eliminating the associated problems, such as the timing problems caused by channel conductors of different length. Additionally, memory module 200 provides for two separate channels 206 and 208 that extend across the memory module in parallel with one another. These two channels provide increased bandwidth for communicating data and other information between the various memory devices 204. Each memory device 204 is coupled to its associated channel such that each of the pins or contacts of the memory device make an electrical connection with a corresponding hole or pad on a substrate 202, such as a printed circuit board. In a particular embodiment, memory devices 204 are surface mount components coupled to both outer surfaces of substrate 202 (i.e., the surface that can be seen in Fig. 2 and the surface on the opposite side of the substrate).

Channels 206 and 208 receive data as indicated in Fig. 2, propagate that data along the channel and past the two memory devices 204. The data is then propagated out the top portion of memory module 200 (as oriented in Fig. 2). As shown in Fig. 2, the channel conductors extend substantially linearly from one edge of substrate 202 to the opposite edge of the substrate. In one embodiment, the top and bottom edges of substrate 200 have electrical contacts or pads (not shown) that allow the channel conductors to communicate with another device, such as a connector designed to receive the edge of the substrate.

In a particular embodiment, channels 206 and 208 are coupled to memory devices 204 located on both sides of the substrate 202. For example, four additional memory devices may be located opposite the four memory devices 204(1) – 204(4) shown on one side of substrate 202. This configuration allows four memory devices 204 to be coupled to each channel 206 and 208. In another embodiment, two additional channels and corresponding memory devices are located on the opposite side of substrate 202. In this embodiment, four channels extend across substrate 202.

Fig. 3 illustrates a perspective view of a memory architecture 250 including a pair of memory modules and a pair of connectors to couple the memory modules to one another. The memory architecture 250 includes a connector 254 mounted to a substrate 252, such as a motherboard or other printed circuit board. A pair of memory modules 256 and 258 each engage slots in connector 254. Each memory module 256, 258 includes multiple memory devices 264 mounted on one or both sides of the memory module. Another connector 260 is coupled, via two slots, to the two memory modules 256 and 258. Connector 260 is also mounted to a

1 substrate 262, which includes electrical conductors (e.g., traces) which
2 communicate signals between the two memory modules 256 and 258.

3 Fig. 4 illustrates an exploded view of the memory architecture shown in
4 Fig. 3. A pair of memory modules 302 and 304 each contain multiple memory
5 devices 306 and 308, respectively. In one embodiment, memory modules 302 and
6 304 are similar to memory module 200 described above with respect to Fig. 2. In
7 this embodiment, each memory module 302 and 304 contains eight memory
8 devices 306 or 308 (i.e., four memory devices on each side of the memory
9 module). The memory architecture contains two channels, each of which follows
10 a channel path 350.

11 A connector 310 is configured to receive the two memory modules 302 and
12 304 such that the electrical contacts (not shown) on the lower edge of each
13 memory module engage a corresponding set of electrical contacts or pads (not
14 shown) in the connector 310. Memory module 302 is inserted into a slot 314 in
15 connector 310 and memory module 304 is inserted into a slot 316 in connector
16 310. The electrical contacts in connector 310 are positioned along the inner
17 surface of slots 314 and 316, as discussed below with reference to Fig. 5. In a
18 particular embodiment, memory modules 302 and 304 are interchangeable.

19 Connector 310 includes multiple conductive pins 312 extending from the
20 bottom of the connector (i.e., the side of the connector opposite the slots 314 and
21 316). The conductive pins 312 engage multiple corresponding apertures (or
22 through-holes) 320 in a substrate 318, such as a printed circuit board. In a
23 particular embodiment, substrate 318 is a motherboard containing a central
24 processing unit (CPU), memory controller, and other components necessary to
25 provide the desired computing functions. Additionally, the motherboard has

1 multiple connectors 310 for receiving multiple pairs of memory modules 302 and
2 304. Depending on the memory storage required, some or all of the connectors
3 310 are populated with memory modules. Memory capacity can be increased in
4 the future by inserting one or more additional pairs of memory modules in unused
5 connectors 310.

6 Another connector 330 engages the top edge of each memory module 302
7 and 304. As discussed above with respect to Fig. 2, the top edge of each memory
8 module also includes multiple electrical contacts or pads. These electrical contacts
9 engage corresponding electrical contacts or pads (not shown) in connector 330.
10 Connector 330 includes a slot 334 to receive memory module 302 and another slot
11 336 to receive memory module 304. The electrical contacts in connector 330 are
12 located along the inner surface of slots 334 and 336, as discussed below with
13 reference to Fig. 5.

14 Connector 330 also includes multiple conductive pins 332 extending from
15 the connector opposite the openings of slots 334 and 336. The conductive pins
16 332 engage multiple corresponding apertures (or through-holes) 340 in a substrate
17 338, such as a printed circuit board. In one embodiment, substrate 338 has a
18 footprint approximately the same size as connector 330. The substrate 338
19 includes multiple conductors for connecting the channel conductors on memory
20 module 302 to corresponding channel conductors on module 304. The multiple
21 conductors are arranged such that all conductors for a particular channel are
22 approximately the same length.

23 As discussed with respect to Fig. 2, each memory module 302 and 304 may
24 have two channels extending across the memory module, each channel terminating
25 at opposite edges of the memory module. Connector 330 and substrate 338

1 provide a mechanism for communicating signals between the top edge conductors
2 on each of the memory modules 302 and 304, thereby allowing each of the
3 channels to extend across both memory modules 302 and 304, and back to the
4 substrate 318 (e.g., the motherboard).

5 For example, broken line 350 (Fig. 4) shows a possible channel path 350.
6 In this example, the channel data is received from substrate 318, through
7 connector 310, and communicated to memory module 304 using the electrical
8 contacts in slot 316 and along the edge of the memory module. The channel
9 continues vertically upward along memory module 304 to connector 330, where
10 the data signals are communicated to the substrate 338. The multiple conductors
11 on (or in) substrate 338 extend the channel across the substrate 338 to the slot 334
12 which provides a connection to memory module 302. The channel then continues
13 vertically downward along memory module 302 to connector 310, where the data
14 signals are communicated back to substrate 318. Thus, a single channel is
15 established by the communication of connector 310, memory modules 302 and
16 304, connector 330 and substrate 338. These separate portions may be referred to
17 as "channel portions" which are combined to create a channel. Any number of
18 channel portions can be coupled to one another to create a channel.

19 In a particular embodiment, the memory architecture 300 includes four
20 channels, all of which follow the same path 350. In this embodiment, two
21 channels follow the outer sides of memory modules 302 and 304 (i.e.,
22 communicating with memory devices 308(4), 308(3), 306(1), and 306(2)). The
23 other two channels follow the inner sides of memory modules 302 and 304 (i.e.,
24 communicating with memory devices 308(2), 308(1), 306(3), and 306(4)). In
25

1 alternate embodiments, channels may alternate between inner and outer sides of the
2 memory modules, thereby providing two channels.

3 As described above, each channel "flows" from right to left (i.e., beginning
4 at memory module 304 and ending at module 302). However, it will be apparent
5 to those skilled in the art that one or more channels can flow in the opposite
6 direction (i.e., beginning at memory module 302 and ending at memory module
7 304). Additionally, examples described herein contain a particular number of
8 memory devices (e.g., four) on each side of a memory module. However, alternate
9 embodiments may contain any number of memory devices on each side of a
10 memory module. Further, each side of a memory module may contain a different
11 number of memory modules and/or a different number of channels. A particular
12 embodiment may contain memory devices only on one side of the memory
13 module. Another embodiment may contain one or more channels only on one side
14 of the memory module. A particular channel may be coupled to any number of
15 memory devices arranged on any number of different memory modules.

16 The embodiment of Fig. 4 illustrates connectors 310 and 330 as having pins
17 that engage through holes in the corresponding substrate 318 or 338. In an
18 alternate embodiment, connectors 310 and 330 are surface mount components
19 mounted to and making electrical contact with the surface of substrate 318 or 338.
20 In other embodiments, any mounting system can be used to mount connectors 310
21 and 330 to corresponding substrates 318 and 338.

22 Fig. 5 is a top view of a connector 400 used to couple two memory modules
23 to one another. Connector 400 corresponds to connectors 310 and 330 discussed
24 above with respect to Fig. 4. Connector 400 includes a body portion 402, typically
25 made of plastic or another insulating material. Connector 400 includes two slots

404 and 406 to receive memory modules or other devices having an edge (e.g., the edge of a substrate) with dimensions and electrical contacts that correspond to the dimensions and electrical contacts of the slots. Each slot 404 and 406 includes multiple electrical contacts or pads 408 along the elongated inner edges of the slots. The size, shape, and spacing of electrical contacts 408 are shown as an example only. Alternate connectors may have any number of electrical contacts of any size and shape, and arranged in any manner.

Connector 400 also includes a pair of openings 410 and 412 in the body 402 of the connector. These openings reduce the amount of material required to manufacture the body 402 and provides visibility through the connector 400 to the attached substrate (not shown). Alternate embodiments may omit one or both of the openings 410 and 412. The same type of connector 400 is used on both edges of a pair of memory modules. Since both slots 404 and 406 are molded or otherwise formed into the same connector, the alignment of the two memory modules with the two slots should be in proper registration. This single connector 400 with two slots 404 and 406 eliminates the need to provide two separate connectors (each with one slot) and align the two connectors on the substrate such that the memory modules are in proper alignment.

Fig. 6 is a flow diagram illustrating a procedure 450 for creating and using a memory architecture of the type described herein. Multiple memory devices (e.g., RDRAMs) are arranged on a substrate such that multiple channels extend from one edge of the substrate to the opposite edge of the substrate (block 452). Typically, the substrate is rectangular (i.e., having two longer (or elongated) sides or edges). In a particular embodiment, the multiple channels extend from one elongated edge to the opposite elongated edge, as shown in Fig. 2. The channel

1 conductors are arranged such that the length of the conductors between opposite
2 edges of the substrate is approximately equal (block 454). Further, the channel
3 conductors extend substantially linearly from one edge of the substrate to the
4 opposite edge of the substrate.

5 A pair of substrates are coupled together using a connector so that each
6 channel extends across a first substrate, across the connector, and across a second
7 substrate (block 456). The channel conductors are arranged on the connector such
8 that all conductors for a particular channel are approximately the same length
9 (block 458). Signals are propagated through the multiple channels to perform
10 memory read, memory write, and memory control operations (block 460).

11 Fig. 7 illustrates another embodiment of a memory architecture 500 having
12 multiple channels in which all channel conductors are substantially the same
13 length. Memory architecture 500 includes four channels that extend across each
14 substrate in parallel with one another. The four channels increase the overall
15 bandwidth available for reading data, storing data, and performing various
16 memory control operations. A first memory module 502 contains eight memory
17 devices 512 on each side of the memory module. Memory module 502 can
18 receive data (including instructions) on four different channels 514, 516, 518, and
19 520. In an alternate embodiment, an additional four channels (not shown) are
20 located on the opposite side of memory module 502.

21 Memory module 502 is coupled to a first slot 504 in a connector. The
22 connector is coupled to a substrate 506, such as a printed circuit board. A second
23 slot 506 in the connector receives a second memory module 510. Memory module
24 510 includes multiple memory devices 522. As shown in Fig. 7, the four channels
25 514-520 extend across the first memory module 502, across substrate 506, and

1 across the second memory module 510. In this arrangement, the conductors of
2 each channel are substantially the same length. Further, the channel conductors
3 follow a substantially linear path across each substrate 502, 506, and 510. A
4 particular channel may have any number of memory devices and have conductors
5 of any length.

6 Thus, a system has been described that provides a multi-channel memory
7 architecture having channel conductors of substantially equal length. Although
8 the description above uses language that is specific to structural features and/or
9 methodological acts, it is to be understood that the invention defined in the
10 appended claims is not limited to the specific features or acts described. Rather,
11 the specific features and acts are disclosed as exemplary forms of implementing
12 the invention.